HPSS@NERSC

BERKELEY LAB

Advancing Data Access to Accelerate Scientific Discovery

NERSC has long been at the forefront of providing its user community with state-of-the-art archival storage since HPSS was deployed at the center in 1998. Today, NERSC's HPSS provides users with a archival storage system offering a theoretical capacity of 8.8 petabytes. This capacity is an example of meeting future needs, as the amount of stored data nearly doubles every year.

In practice, NERSC's HPSS currently houses 2 petabytes of data stored in 53 million files. On any given day, between 4 and 8 terabytes of data are transferred in and out of the system. The systems theoretical maximum throughput is 4.2 gigabytes per second.

NERSC is one of the original HPSS development partners, along with LLNL, LANL, SNL, ORNL, and IBM. NERSC is the only site where this effort is funded by Office of Science funds, yet HPSS is critical to a number of SC projects that are independent of NERSC. Further, HPSS is the only archive storage system that can meet the scalability, performance, and feature requirements for capability computing. NERSC will continue to participate in the HPSS collaboration, which brings a number of benefits, including early access to new features and new technology, ability to prioritize requirements and strategy, and collaboration with NNSA activities.

In addition to providing large-scale storage resources to scientific researchers, NERSC's HPSS team is also collaborating with developers at other sites to make HPSS an even more accessible and valuable tool for computational science.

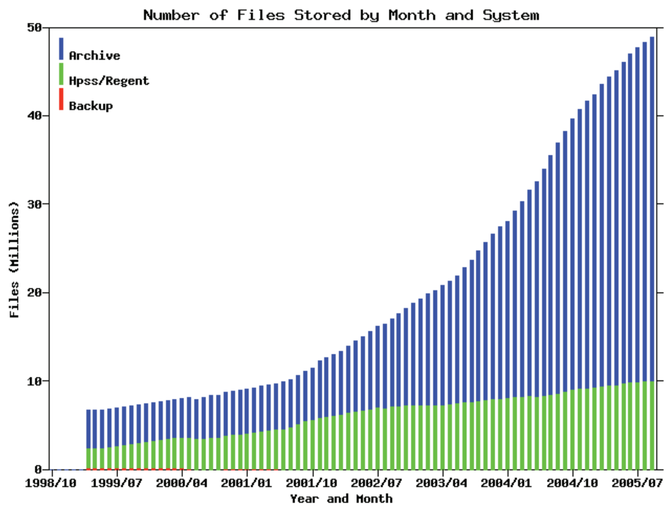

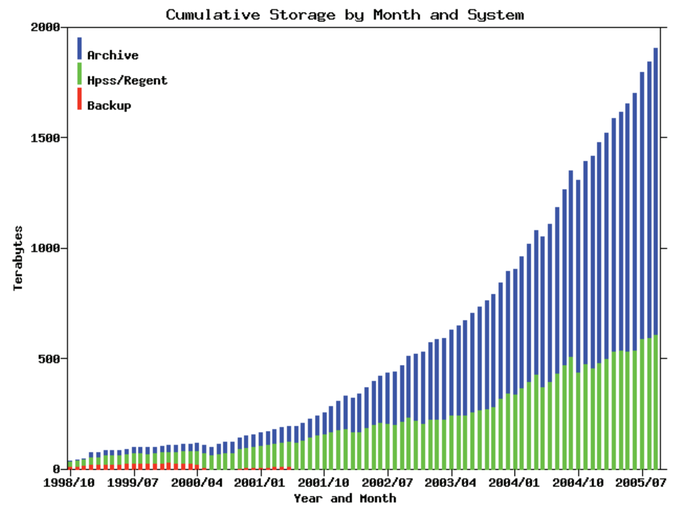

Large Repositories for Scientific Data

From global climate change to particle physics, NERSC's HPSS supports a wide range of scientific projects by researchers at national laboratories and universities around the country. Here are examples of some of the largest users of the system.

STAR DETECTOR SIMULATIONS AND DATA ANALYSIS
Brookhaven National Laboratory's Solenoid Tracker at RHIC (STAR) generates millions of particle collisions per experiment, but only a handful of such collisions will provide clues in the hunt for the quark gluon plasma, the theorized atomic outcome of the Big Bang. Storing, retrieving and analyzing the data make STAR a significant user of NERSC's HPSS. Currently, some 77.5 million files account for 2.8 terabytes of data in the system.

DATA ANALYSIS AND SIMULATIONS FOR KAMLAND
The KamLAND experiment, located in a mine in Japan, is a 1 kiloton liquid scintillator detector that was built to study anti-neutrinos coming from Japanese nuclear reactors, KamLAND records data 24 hours a day, seven days a week. These data are shipped on tapes from the experimental site to LBNL, where they are read off the tapes and stored in the HPSS at NERSC. KamLAND records about 200 GB of data each day and HPSS currently has more than 250 TB of KamLAND data stored in 10.6 million files.

Analysis of the KamLAND data led to the first observation of geologically produced anti-neutrinos. The findings were featured on the cover of the July 28, 2005 issue of Nature.

THE NEARBY SUPERNOVA FACTORY
In their quest to learn more about dark energy, scientists need more more observations of Type Ia supernovae. The Nearby Supernova Factory (SNfactory) led by Berkeley Lab. Was established to reduce the uncertainty and improve the calibration of Type Ia supernovae. To do this, the SNfactory needs to discover and make detailed observations of 300 to 600 low-redshift supernovae, many more than have been studied so far. During its first year of operation, the SNfactory found 34 supernovae, the best performance ever for a "rookie" supernova search. The group is now discovering supernovae at the rate of nine a month.

The Jet Propulsion Laboratory's Near Earth Asteroid Tracking program (NEAT) uses remote telescopes in Southern California and Hawaii to record 500 square degrees of the sky to NERSC each night, amounting to 50 gigabytes. The data pipeline software automatically archives these in NERSC's HPSS. NEAT's telescopes revisit the same regions of the sky roughly every six days during a typical 18-day observing period. When a supernova appears in one of those galaxies, the SNfactory can find it using image subtraction software that can sift through billions of objects. This analysis is done using NERSC's PDSF Linux cluster. Currently, the SNFactory has about 177,000 files containing 690 gigabytes of data.

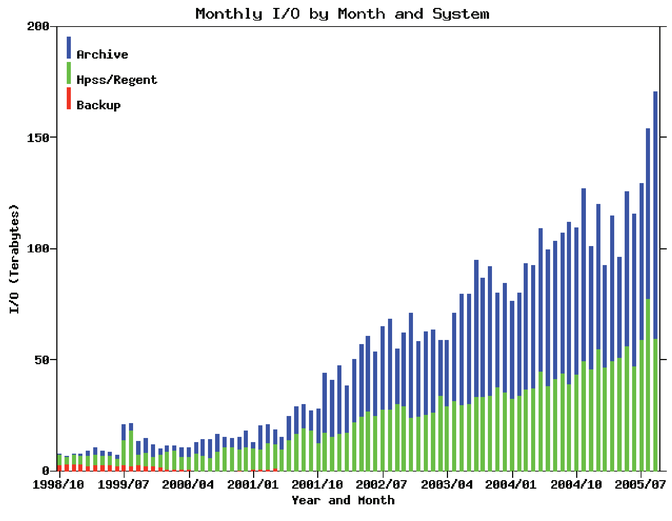

Current Development Collaborations

NERSC's Mass Storage Group is currently involved in two significant development collaborations aimed at advancing the usefulness of HPSS both at NERSC and at other sites. Capabilities from both collaborations are slated to be included in the 2006release of HPSS 6.2.

INTEGRATING HPSS INTO GRIDS
Argonne National Laboratory, with its expertise in all things Grid, is collaborating with IBM and tapping NERSC's HPSS expertise to integrate HPSS accessibility into the Globus Toolkit for Grid applications. NERSC is providing the testbed system for this collaboration

At Argonne, researchers are adding functionality to the Grid Daemon so that the appropriate class of service can be requested from HPSS. IBM is contributing the development of an easy-to-call library of parallel I/O routines that work with HPSS structures and are also easy to integrate into the gridftp deamon. This library will ensure that Grid FTP requests to HPSS movers are handled correctly.

NERSC is providing the HPSS platform and testbed system for IBM and ANL to do their respective development projects. As pieces are completed, NERSC tests the components and works with the developers to help identify and resolve problems.

The public release of this capability is scheduled with HPSS 6.2 in spring 2006, as well as future releases of the Globus Toolkit.

GPFS/HPSS DEVELOPMENT
The Mass Storage Group is collaborating with IBM and the San Diego Supercomputer Center to develop a Hierarchical Storage Manager that can be used with IBM's Global Parallel File System (GPFS). The HSM capability with GPFS will provide a recoverable GPFS File System that is transparent to users and fully backed up and recoverable from NERSC's multi-petabyte archive on HPSS.

One of the key capabilities of the HSM is that users' files will automatically be backed up on HPSS as they are created. Additionally, files on the GPFS which have not been accessed for a specified period of time will be automatically purged from the system as space is needed by users for files currently in use. Since the purged files were already migrated/backed up on HPSS, they can easily by recovered by users when needed. HSM will stage the purged files automatically.

The benefits of this approach will lead t:

- An infinitely large file system: Since files will be automatically backed up, purged and retrievable thanks to NERSC's multi-petabyte archival storage, there will virtually be no limits on the size and number of files in the system.
- A better computing environment for users. Instead of needing to use utilities to move files, users will work in a transparent system – they won't need to know where there files are stored in order to access them. There will also be greater reliability and files will be able to be restored if needed.

The first system test of HSM will be at the San Diego Supercomputer Center and is scheduled to begin in December 2005. In spring of 2006, HSM will begin testing on Linux systems at NERSC/

The deployment of the GPFS/HPSS HSM will be on the Facility-Wide File System (FWFS), which NERSC is deploying. The FWFS will provide consolidated storage for online user data at NERSC, replacing traditional system-local parallel file systems for home directories, scratch storage, and project storage. FWFS will grow and evolve over time, serving several generations of computational systems. By providing a single unified namespace, FWFS will make it easier for users to manage their data across multiple systems. They will no longer need to copy data between NERSC systems for pre- and post-processing.

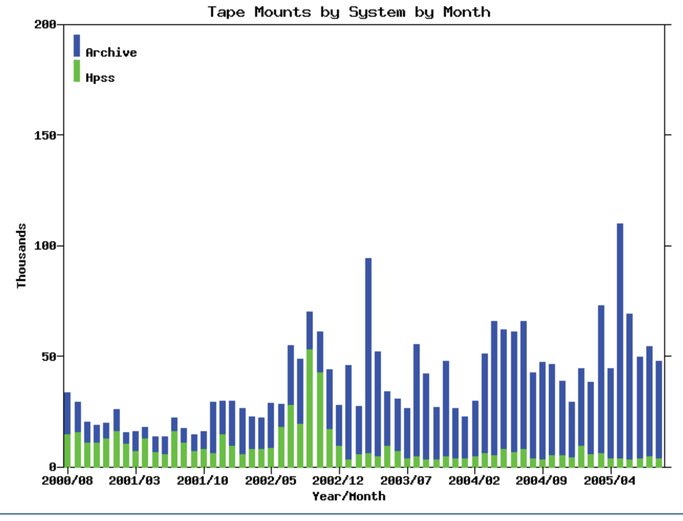

Office of Science
U.S. DEPARTMENT OF ENERGY

NERSC
NATIONAL ENERGY RESEARCH SCIENTIFIC COMPUTING CENTER